

Please type a plus sign (+) inside this box → ☐

Approved for use through 09/30/2000. OMB 0651-0032
Patent and Trademark Office: U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it displays a valid OMB control number.

UTILITY PATENT APPLICATION TRANSMITTAL <i>(Only for new nonprovisional applications under 37 C.F.R. § 1.53(b))</i>	Attorney Docket No.	1662-28800 (P00-2998)
	First Inventor or Application Identifier	William P. BUNTON et al.
	Title	High-Speed Interconnection Link Having Automated Crossed ...
	Express Mail Label No.	EL497841418US

APPLICATION ELEMENTS See MPEP chapter 600 concerning utility patent application contents	ADDRESS TO: Assistant Commissioner for Patents Box Patent Application Washington, DC 20231
1. <input checked="" type="checkbox"/> * Fee Transmittal Form (e.g., PTO/SB/17) <i>(Submit an original and a duplicate for fee processing)</i> 2. <input checked="" type="checkbox"/> Specification [Total Pages 23] <i>(preferred arrangement set forth below)</i> - Descriptive title of the Invention (plus cover sheet) - Cross References to Related Applications - Statement Regarding Fed sponsored R & D - Reference to Microfiche Appendix - Background of the Invention - Brief Summary of the Invention - Brief Description of the Drawings (if filed) - Detailed Description - Claim(s) - Abstract of the Disclosure 3. <input checked="" type="checkbox"/> Drawing(s) (35 U.S.C. 113) [Total Sheets 7] 4. Oath or Declaration [Total Pages 1] a. <input checked="" type="checkbox"/> Newly executed (original or copy) b. <input type="checkbox"/> Copy from a prior application (37 C.F.R. § 1.63(d)) <i>(for continuation/divisional with Box 16 completed)</i> i. <input type="checkbox"/> DELETION OF INVENTOR(S) Signed statement attached deleting inventor(s) named in the prior application, see 37 C.F.R. §§ 1.63(d)(2) and 1.33(b)	5. <input type="checkbox"/> Microfiche Computer Program (Appendix) 6. Nucleotide and/or Amino Acid Sequence Submission <i>(if applicable, all necessary)</i> a. <input type="checkbox"/> Computer Readable Copy b. <input type="checkbox"/> Paper Copy (identical to computer copy) c. <input type="checkbox"/> Statement verifying identity of above copies
ACCOMPANYING APPLICATION PARTS	
7. <input checked="" type="checkbox"/> Assignment Papers (cover sheet & document(s)) 8. <input type="checkbox"/> 37 C.F.R. § 3.73(b) Statement of Power of Attorney <i>(when there is an assignee)</i> 9. <input type="checkbox"/> English Translation Document (if applicable) 10. <input type="checkbox"/> Information Disclosure Statement (IDS)/PTO-1449 [Copies of IDS Citations] 11. <input type="checkbox"/> Preliminary Amendment 12. <input checked="" type="checkbox"/> Return Receipt Postcard (MPEP 503) <i>(Should be specifically itemized)</i> 13. <input type="checkbox"/> * Small Entity Statement filed in prior application, Status still proper and desired <i>(PTO/SB/09-12)</i> 14. <input type="checkbox"/> Certified Copy of Priority Document(s) <i>(if foreign priority is claimed)</i> 15. <input type="checkbox"/> Other:	

*** NOTE FOR ITEMS 1 & 13: IN ORDER TO BE ENTITLED TO PAY SMALL ENTITY FEES, A SMALL ENTITY STATEMENT IS REQUIRED (37 C.F.R. § 1.27), EXCEPT IF ONE FILED IN A PRIOR APPLICATION IS RELIED UPON (37 C.F.R. § 1.28).**

16. If a CONTINUING APPLICATION, check appropriate box, and supply the requisite information below and in a preliminary amendment:
☐ Continuation ☐ Divisional ☐ Continuation-in-part (CIP) of prior application No. _____
 Prior application information: Examiner _____ Group / Art Unit _____

For CONTINUATION or DIVISIONAL APPS only: The entire disclosure of the prior application, from which an oath or declaration is supplied under Box 4b, is considered a part of the disclosure of the accompanying continuation or divisional application and is hereby incorporated by reference. The incorporation can only be relied upon when a portion has been inadvertently omitted from the submitted application parts.

17. CORRESPONDENCE ADDRESS					
<input type="checkbox"/> Customer Number or Bar Code Label		23505 <i>(Insert Customer No. or Attach bar code label here)</i>		or <input checked="" type="checkbox"/> Correspondence address below	
Name	Jonathan M. Harris				
	Conley, Rose & Tayon, P.C.				
Address	PO Box 3267				
City	Houston	State	TX	Zip Code	77253-3267
Country	USA	Telephone	713-238-8000	Fax	713-238-8008

Name (Print/Type)	Michael F. Heim	Registration No. (Attorney/Agent)	32,702
Signature	<i>Michael F. Heim</i>	Date	June 20, 2000

Burden Hour Statement. This form is estimated to take 0.2 hours to complete. Time will vary depending upon the needs of the individual case. Any comments on the amount of time you are required to complete this form should be sent to the Chief Information Officer, Patent and Trademark Office, Washington, DC 20231. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. SEND TO: Assistant Commissioner for Patents, Box Patent Application, Washington, DC 20231.

M:\C\1662\28800\PTO TRANS 01

JC836 U.S. PTO
09/597192
06/20/00

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE
APPLICATION FOR UNITED STATES LETTERS PATENT

**HIGH-SPEED INTERCONNECTION ADAPTER HAVING
AUTOMATED CROSSED DIFFERENTIAL PAIR CORRECTION**

By:

William Bunton
415 Greenway Dr.
Pflugerville, Texas 78660
Citizenship: U.S.A.

John Krause
1310 E. University Ave.
Georgetown, Texas 78626
Citizenship: U.S.A.

HIGH-SPEED INTERCONNECTION ADAPTER HAVING AUTOMATED CROSSED DIFFERENTIAL PAIR CORRECTION

5 CROSS-REFERENCE TO RELATED APPLICATIONS

Not applicable.

STATEMENT REGARDING FEDERALLY SPONSORED RESEARCH OR DEVELOPMENT

Not applicable.

10 BACKGROUND OF THE INVENTION

Field of the Invention

The present invention generally relates to high bandwidth interconnections for use in networking environments such as local area networks (LAN), wide area networks (WAN) and storage area networks (SAN). More specifically, it relates to a method of correcting inverted
15 binary signals arising from crossing of differential pair signal carriers.

Description of Related Art

Internet and electronic commerce has grown to the point where demands placed on existing computer systems are severely testing the limits of system capacities. Microprocessor and
20 peripheral device performances have improved to keep pace with emerging business and educational needs. For example, internal clock frequencies of microprocessors have increased dramatically, from less than 100 MHz to more than 1 GHz over a span of less than ten years. Where this performance increase is inadequate, high performance systems have been designed with multiple processors and clustered architecture. It is now commonplace for data and software
25 applications to be distributed across clustered servers and separate networks. The demands created

by these growing networks and increasing speeds are straining the capabilities of existing Input/Output (I/O) architecture.

Peripheral Component Interconnect (PCI), released in 1992, is perhaps the most widely used I/O technology today. PCI is a shared bus-based I/O architecture and is commonly applied as a means of coupling a host computer bus (front side bus) to various peripheral devices in the system. Publications that describe the PCI bus include the *PCI Specification, Rev. 2.2*, and *Power Management Specification 1.1*, all published by the PCI Special Interest Group. The principles taught in these documents are well known to those of ordinary skill in the art and are hereby incorporated herein by reference.

At the time of its inception, the total raw bandwidth of 133 MBps (32 bit, 33 MHz) provided by PCI was more than sufficient to sustain the existing hardware. Today, in addition to microprocessor and peripheral advancements, other I/O architectures such as Gigabit Ethernet, Fibre Channel, and Ultra3 SCSI are outperforming the PCI bus. Front side buses, which connect computer microprocessors to memory, are approaching 1-2 GBps bandwidths. It is apparent that the conventional PCI bus architecture is not keeping pace with the improvements of the surrounding hardware. The PCI bus is quickly becoming the bottleneck in computer networks.

In an effort to meet the increasing needs for I/O interconnect performance, a special workgroup led by Compaq Computer Corporation developed PCI-X as an enhancement over PCI. The PCI-X protocol enables 64-bit, 133 MHz performance for a total raw bandwidth that exceeds 1 GBps. While this is indeed an improvement over the existing PCI standard, it is expected that the PCI-X bus architecture will only satisfy I/O performance demands for another two or three years.

In addition to the sheer bandwidth limitations of the PCI bus, the shared parallel bus architecture used in PCI creates other limitations which affect its performance. Since the PCI bus

is shared, there is a constant battle for resources between processors, memory, and peripheral devices. Devices must gain control of the PCI bus before any data transfer to and from that device can occur. Furthermore, to maintain signal integrity on a shared bus, bus lengths and clock rates must be kept down. Both of these requirements are counter to the fact that microprocessor speeds
5 are going up and more and more peripheral components are being added to today's computer systems and networks.

Today, system vendors are decreasing distances between processors, memory controllers and memory to allow for increasing clock speeds on front end buses. The resulting microprocessor-memory complex is becoming an island unto itself. At the same time, there is a
10 trend to move the huge amounts of data used in today's business place to storage locations external to network computers and servers. This segregation between processors and data storage has necessitated a transition to external I/O solutions.

One solution to this I/O problem has been proposed by the Infiniband(SM) Trade Association. The Infiniband(SM) Trade Association is an independent industry body that is
15 developing a channel-based, switched-network-topology interconnect standard. This standard will de-couple the I/O subsystem from the microprocessor-memory complex by using I/O engines referred to as channels. These channels implement switched, point to point serial connections rather than the shared, load and store architecture used in parallel bus PCI connections.

The Infiniband interconnect standard offers several advantages. First, it uses a differential
20 pair of serial signal carriers, which drastically reduces conductor count. Second, it has a switched topology that permits many more nodes which can be placed farther apart than a parallel bus. Since more nodes can be added, the interconnect network becomes more scalable than the parallel bus network. Furthermore, as new devices are added, the links connecting devices will fully

support additional bandwidth. This Infiniband architecture will let network managers buy network systems in pieces, linking components together using long serial cables. As demands grow, the system can grow with those needs.

The trend towards using serial interconnections as a feasible solution to external I/O solutions is further evidenced by the emergence of the IEEE 1394 bus and Universal Serial Bus (USB) standards. USB ports, which allow users to add peripherals ranging from keyboards to biometrics units, have become a common feature in desktop and portable computer systems. USB is currently capable of up to 12 MBps bandwidths, while the IEEE 1394 bus is capable of up to 400 MBps speeds. A new version of the IEEE 1394 bus (IEEE 1394b) can support bandwidth in excess of 1 GBps.

Maintaining signal integrity is extremely important to minimize bit error rates (BER). At these kinds of bandwidths and transmission speeds, a host of complications which affect signal integrity may arise in the physical layer of a network protocol. The physical layer of a network protocol involves the actual media used to transmit the digital signals. For Infiniband, the physical media may be a twisted pair copper cable, a fiber optic cable, or a copper backplane. Interconnections using copper often carry the transmitted signals on one or more pairs of conductors or traces on a printed circuit board. Each optical fiber or differential conductor pair is hereafter called a “lane”.

Where multiple lanes are used to transmit serial binary signals, examples of potential problems include the reordering of the lanes and skew. Skew results from different lane lengths or impedances. Skew must be corrected so data that is transmitted at the same time across several lanes will arrive at the receiver at the same time. Lane reordering must be corrected so a digital signal may be reconstructed and decoded correctly at the receiver end.

Even in the simplest case involving a single differential wire pair, a potential problem exists in the routing of the differential wire pair. It is possible for wires to be crossed either inadvertently, as in a cable miswire, or intentionally, as may be necessary to minimize skew. In transmitting digital signals via a differential wire pair, one wire serves as a reference signal while the other wire transmits the binary signal. If the wire terminations are incorrect, the binary signal will be inverted.

Conventional correction and prevention of these types of problems has been implemented with the meticulous planning and design of signal paths. Differential wire pairs are typically incorporated into cables as twisted wire pairs of equal lengths. However, matched delay or matched length cabling is more expensive because of the manufacturing precision required. In backplane designs, trace lengths may vary because of board congestion, wire terminations and connector geometries. Shorter traces are often lengthened using intentional meandering when possible to correct for delay caused by other components. It is often impractical to correct crossed differential pairs because one trace passes through two vias to “cross under” the other trace. The vias introduce a substantial time delay, thereby causing data skew. Alternatively, the differential pairs are left uncorrected and the data inversion is accounted for using pin straps or boundary scan techniques. Both of these fixes require intervention by the system designer. These techniques have also been used to correct lane reversal.

The physical layer in Infiniband carries signals encoded by a digital transmission code called “8B/10B”. 8B/10B is an encoding/decoding scheme which converts an 8-bit word (i.e., a byte) at the link layer of the transport protocol to a 10-bit word that is transmitted in the physical layer of the same protocol.

The 8B/10B code is a “zero-DC” code, which provides some advantages for fiber optic and copper wire links. Transmitter level, receiver gain, and equalization are simplified and their precision is improved if the signals have a constant average power and no DC component. Simply stated, in converting an 8-bit word to a 10-bit word, the encoder selects the 10-bit representation based on the sign of the running disparity of the digital signal. Running disparity refers to a running tally of the difference between the number of 1 and 0 bits in a binary sequence. If the running disparity is negative (implying that more 0 bits have been transmitted than 1 bits), the subsequent 8B/10B word will contain more 1 bits than 0 bits to compensate for the negative running disparity. In the 8B/10B code, every 8-bit word has two 10-bit equivalent words. The 10-bit equivalent words will have five or more 1 bits for a negative running disparity and five or more 0 bits for a positive running disparity. For a more detailed description of the 8B10B code, refer to Widmer and Franaszek, “A DC-Balanced, Partitioned-Block, 8B/10B Transmission Code”, IBM J. Res. Develop., Vol. 27, No. 5, September 1983, which is hereby incorporated by reference.

The above design considerations clearly make physical layer (i.e., cables, backplanes) manufacturing a difficult venture in high clock frequency systems. Design costs and manufacturing costs are drastically increased due to the need to alleviate these types of problems. It is desirable, therefore, to provide a method of automatically correcting these types of errors with information embedded in the signals. It is further desirable to develop a method of automatically detecting and correcting inverted signals resulting from crossed differential signal pairs. This method of correction may advantageously allow for less stringent design requirements and could decrease design and manufacturing costs for high bandwidth data links.

BRIEF SUMMARY OF THE INVENTION

The problems noted above are solved in large part by an adapter configured to automatically detect and compensate for differential signal inversion. In one embodiment, the adapter is part of a computer network having differential conductor pairs conveying differential signals between network devices. The network devices include adapters coupled to transmit and receive signals via the differential conductor pairs. The adapter preferably includes a lane receiver, a decoder, and a synchronization circuit. The lane receiver is configured to receive a single differential signal and to convert the differential signal into a sequence of code symbols. The decoder decodes the code symbols to produce a sequence of received symbols. The synchronization circuit examines the sequence of received symbols to determine if it is incorrect due to inversion of the differential signal, and if so, it causes the lane receiver to correct for the differential signal inversion. It is expected that the received symbol sequence will include a training symbol sequence which will have a start symbol whose decoded value is unaffected by differential signal inversion, and a training symbol whose decoded value is indicative of the presence or absence of inversion. The synchronization circuit, upon identifying the training sequence, will thus be able to determine whether inversion exists and be able to automatically correct for it.

BRIEF DESCRIPTION OF THE DRAWINGS

For a detailed description of the preferred embodiments of the invention, reference will now be made to the accompanying drawings in which:

Figure 1 shows an illustrative diagram of a simple computer network which supports serial I/O connections;

Figure 2 shows a functional block diagram of a simple computer network which supports serial I/O connections;

Figure 3 shows a functional block diagram of an alternative computer network which supports serial I/O connections;

5 Figure 4 shows a ladder diagram of the training sequence used to train ports that are coupled to opposite ends of a serial physical link;

Figure 5 shows a table of the preferred training packets used to train ports that are coupled to opposite ends of a serial physical link;

10 Figure 6 shows a table of the preferred lane identifiers used to label the individual channels in a serial physical link; and

Figure 7 shows a functional block diagram of a serial physical link.

NOTATION AND NOMENCLATURE

15 Certain terms are used throughout the following description and claims to refer to particular system components. As one skilled in the art will appreciate, computer companies may refer to a component by different names. This document does not intend to distinguish between components that differ in name but not function. In the following discussion and in the claims, the terms “including” and “comprising” are used in an open-ended fashion, and thus should be interpreted to mean “including, but not limited to...”. Also, the term “couple” or “couples” is intended to mean
20 either an indirect or direct electrical connection. Thus, if a first device couples to a second device, that connection may be through a direct electrical connection, or through an indirect electrical connection via other devices and connections.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENTS

Figure 1 shows an example of a computer network representing a preferred embodiment, in which a central computer 100 is coupled to an external storage tower 110 and a network router 120 via a multiservice switch 130. Storage tower 110 may be internally connected by a Fibre Channel, SCSI, or any suitable storage network. Network router may be connected to a LAN (local area network) or ISDN (Integrated Services Digital Network) network or it may provide a connection to the internet via a suitable ATM (asynchronous transfer mode) network. It should be appreciated that any number of computers, servers, switches, hubs, routers, or any suitable network device can be coupled to the network shown in Figure 1.

In the preferred embodiment shown in Figure 1, the devices are connected via a point to point serial link 140. The serial link may comprise an even number of lanes or channels through which data is transmitted. Of the even number of lanes, half will transmit serial data in one direction while the other half transmits data in the opposite direction. In the preferred embodiment, the physical links will implement 1, 4, or 12 lanes in each direction. Thus, each link will have a total of 2, 8, or 24 total lanes.

In the latter two implementations (i.e., the 4 and 12 lane links), a single stream of bytes arriving at the input to the physical link are distributed evenly, or “striped”, among the multiple lanes. In the case of the 12-lane link, the first byte is sent to the first lane, the second byte is sent to the second lane and so on until the 12th byte is sent to the 12th lane. At that point, the byte distribution cycles back to the first lane and the process continues. Thus, over time, each lane will carry an equal 1/12th share of the bandwidth that the entire link carries. The same process and technique are used in the 4 lane link. Alternative embodiments with different numbers of lanes would preferably implement this striping process.

Once the bytes are distributed among the individual lanes, the 8-bit words are encoded into 10-bit words and transmitted through the physical link. At the output of the physical link, the 10-bit words are decoded back to 8-bit bytes and are re-ordered to form the original stream of 8-bit words.

5 Figure 2 represents a functional block diagram of the computer network shown in Figure 1. The computer 100 generally includes a central processor unit (CPU) 202, a main memory array 204, and a bridge logic device 206 coupling the CPU 202 to the main memory 204. The bridge logic device is sometimes referred to as a “North bridge” for no other reason than it often is depicted at the upper end of a computer system drawing. The North bridge 206 couples the CPU
10 202 and memory 204 to various peripheral devices in the system through a primary expansion bus (Host Bus) such as a Peripheral Component Interconnect (PCI) bus or some other suitable architecture.

The North bridge logic 206 also may provide an interface to an Accelerated Graphics Port (AGP) bus that supports a graphics controller 208 for driving the video display 210. If the
15 computer system 100 does not include an AGP bus, the graphics controller 208 may reside on the host bus.

Various peripheral devices that implement the host bus protocol may reside on the host bus. For example, a modem 216, and network interface card (NIC) 218 are shown coupled to the host bus in Figure 2. The modem 216 generally allows the computer to communicate with other
20 computers or facsimile machines over a telephone line, an Integrated Services Digital Network (ISDN), or a cable television connection, and the NIC 218 permits communication between computers over a local area network (LAN) (e.g., an Ethernet network card or a Cardbus card).

These components may be integrated into the motherboard or they may be plugged into expansion slots that are connected to the host bus.

Figure 2 also depicts a host channel adapter (HCA) 220 connected to the host bus and target channel adapters (TCA) 230, 240 connected to the external network devices 110, 120.

5 These channel adapters generally provide address and translation capability for the switched topology architecture in the preferred embodiment. The channel adapters 220, 230, 240 preferably have dedicated IPv6 (Internet Protocol Version 6) addresses that can be recognized by the network switch 130. As data is transmitted to the network, the source file is divided into packets of an efficient size for routing. Each of these packets is separately numbered and includes the address of
10 the destination. When the packets have all arrived, they are reassembled into the original file. The network switch 130 in this preferred embodiment can detect the destination address, and route the data to the proper location.

Figure 2 also shows the physical links 140 between the network devices as simple two lane links. In the embodiment shown in Figure 2, data would flow through one lane in one direction
15 while data would flow through the parallel lane the other direction. As discussed above, alternative embodiments comprising any even number of lanes are also permissible, with 2, 8, and 24 lanes per link being the preferred number.

Figure 3 shows an alternative embodiment of the computer network in which the computer
20 100 is replaced by a server 300 with a simple memory-processor architecture. Such a server may be part of a cluster of servers, a group of several servers that share work and may be able to back each other up if one server fails. In this particular embodiment, the server 300 is coupled to the switched-fabric network in much the same way the computer 100 of Figure 1 is connected. The physical link 140 is connected to the server via a host channel adapter (HCA) 220. However, in

this embodiment, the HCA 220 is connected directly to a North Bridge 206. Alternatively, the HCA 220 may be connected directly to a memory controller. In either event, a shared peripheral bus, such as a PCI bus, is not necessary in this embodiment. A peripheral bus may still be used in the server 300, but is preferably not used to couple the north bridge 206 to the HCA 220.

5 As discussed above, the serial data sent through the physical links is sent in the form of packets. The preferred embodiment uses the idea of packetized data and uses specialized packets called Training Set 1 and Training Set 2 to train the network devices prior to transmitting “real” data through the switched network. The actual content and structure of the training sets are discussed in further detail below.

10 Figure 4 shows a link training ladder diagram describing the sequence of events during the training of ports located on either side of the physical link. In the preferred embodiment, a port refers to a transmitting and receiving device configured with a channel adapter to communicate via a serial link. In Figure 4, Port A 400 refers to one such device while Port B 410 refers to the device at the other end of the serial link.

15 The training data, TS1 420 and TS2 430 are packets of known data that are transmitted between Port A 400 and Port B 410. The purpose behind the training sets are twofold. First, the initiation and duration of the training sequence is established by the transmission and reception of the training sets. Secondly, given that the training sets contain pre-determined data, the transmit and receive ports can use this knowledge to correct for any errors (e.g., data inversion, lane skew)
20 that may result during transmission through the physical link. Since the errors are a constant, permanent result of routing in the physical media, the training sequence may be used to automatically correct the errors for all subsequent data transferred through that physical link.

Figure 4 represents a time line for both Port A 400 and Port B 410 with time elapsing toward the bottom of the figure. Before training begins, Port A 400 may exist in an enabled state 440 while Port B is in a disabled or link down state 450. By transmitting an initial sequence of TS1 training sets 420, Port A 400 can effectively wake up Port B 410 from a disabled state to an enabled state 440. Once Port B is enabled 440, two things occur. First, Port B 410 will begin transmitting TS1 training sets back to Port A 400. Secondly, Port B 410 will check the content of the incoming TS1 training sets 420 to see if the data was received as it was sent. If there is any discrepancy, Port B 410 will correct the incoming signals so that the original content of TS1 420 is restored. At this point, Port B 410 will be trained 460 and will respond by sending the second training set, TS2 430, back to Port A 400.

Meanwhile, Port A 400 has been receiving TS1 data 420 from Port B 410 and performs the same signal integrity checks and correction that Port B has completed. Once both ports are trained with TS1 data 420, the ports will proceed by sending TS2 training data 430. This second training set serves as a redundancy check to verify that the Ports were trained properly with TS1 data 420. In addition, the TS2 data 430 signifies that both ports are trained and are ready to transmit and receive data packets 470. Once a port is transmitting and receiving the TS2 training sequence, it may begin sending data. With physical link errors corrected by the training sequences, the data packets 480 can then be transmitted and received by the ports as intended.

In the event the training sequence fails, a timeout may occur and the affected port may be powered down or otherwise deactivated. Thus, when a transmission fault occurs, locating the problems in the physical link is facilitated by determining which port has been deactivated. By comparison, failure isolation in a bus architecture can be difficult because if one attached device

fails, the entire system may fail. Discovering which device caused the failure is typically a hit-or-miss proposition.

Figure 5 shows the actual format and content of the training sets TS1 and TS2. In the preferred embodiment, each training set is 16 words long. It should be appreciated however, that training sets of different lengths are certainly possible. The width of the training set corresponds to the number of physical lanes in a training set. In the preferred embodiment, the training sets are 1, 4, or 12 words wide corresponding to the 1, 4, and 12 lanes in the preferred embodiment of the physical link. Certainly, other combinations of lane quantities are possible, but the width of the training set corresponds to the number of lanes in the physical link. The embodiment shown in

Figure 5 corresponds to a 4 lane link.

Each word in the training set is a 10-bit word that complies with the 8B/10B code discussed above. The first row (COM) in each column is a comma delimiter with a preferred code name K28.5. The second row in each column is a lane identifier that is unique to each lane in the physical link. A table of preferred lane identifiers is shown in Figure 6. In a single lane link, only lane identifier 0 is used. In a 4 lane link, lane identifiers 0, 1, 2, and 3 are used. In a 12 lane link, all twelve lane identifiers shown in Figure 6 are used. After the lane identifier, the remaining 14 rows of the 16 row training sets are repeated 10-bit words. For training set 1, the repeated word name is D10.2. For training set 2, the repeated word name is D5.2.

The comma delimiter and lane identifiers are chosen to be insensitive to data inversion. That is, inverting a comma delimiter or a lane identifier symbol changes only the running disparity and not the symbol itself. Consider the 10-bit word for the comma delimiter K28.5. For a negative running disparity, the word is 001111 1010. For a positive running disparity, the word is 110000 0101. These two words are complements of each other. Inverting all the bits in the first word will

yield the second word and vice-versa. Hence, regardless of whether or not a bit inversion has occurred in the physical link, when the receiver port decodes this word, the comma delimiter will result. The same is also true for each of the lane identifiers in Figure 6. For each lane identifier, the 10-bit words for negative running disparity are the complement of the 10-bit word for positive running disparity. Thus, a receiver will always know when a comma delimiter has arrived and which lane identifier corresponds to a given bit stream. The preferred code names selected for the comma delimiter and the lane identifiers were selected because of their inversion properties. Other code words exhibiting the same properties will also work in alternative embodiments.

To correct for a crossed differential pair (i.e., bit reversal in a lane), a receiver will decode the 10-bit words arriving in a particular bit stream and determine when the training data starts (as marked by the comma delimiter) and determine the lane number (as marked by the lane identifier). For training set 1, the preferred 10-bit code name is D10.2 and the bit sequence for positive running disparity is 010101 0101. The D10.2 code word is chosen for the training set because it uses the exact same code word for negative running disparity as it does for positive running disparity. Thus, the receiver expects to receive the 010101 0101 sequence repeated 14 times for each training set 1 packet regardless of the current state of the running disparity. However, if the complementary code word is received (101010 1010), a completely different word is decoded. The inverted word corresponds to the D21.5 code word. If the receiver decodes this inverse word, the decoder will be configured to invert all the subsequent bits received in that particular lane.

The same conditions hold true for training set number 2. For training set 2, the preferred 10-bit code name is D5.2 and the bit sequence for both positive and negative running disparity is 101001 0101. The inverted code word (010110 1010) corresponds to code name D26.5. Again, the receiving port will attempt to recover the correct polarity by inverting a bit stream if the D26.5

code words are received. The preferred code names selected for training set 1 and training set 2 were selected because of their inversion properties. Other code words exhibiting the same properties will also work in alternative embodiments.

Figure 7 shows a block diagram of a preferred embodiment of a serial physical link.

5 Included in the link are Port A 400 and Port B 410 as discussed above. The link shown in Figure 7 is a 2-lane link with one lane configured to transmit in one direction and the other lane configured to transmit in the opposite direction. Included in the link are retimers 700, 710 located at opposite ends of the link. Retimers 700, 710 provide a means of compensating for minor clock tolerances that result in different clock rates between Port A 400 and Port B 410. To compensate for these
10 clock differences, a data packet called a SKIP ordered set 720 is transmitted at regular intervals amidst the training, data, or idle data packets. In the preferred embodiment, the SKIP ordered sets 720 are inserted every 4608 symbol clocks and include a COM delimiter followed by three SKIP words. As with the training sets, the SKIP ordered sets 720 are as wide as the number of lanes in the link. In Figure 7, the link contains only one lane, so the SKIP ordered sets 720, contain only
15 one column of 10-bit words.

If a delay is needed to compensate for advanced clock timing, the retimers 700, 710 may insert an additional SKIP word to delay the arrival of subsequent data at the receiving end of the link. This scenario is depicted by the SKIP ordered set 740 shown at the receiver of Port B 410. SKIP ordered set 740 includes two additional SKIP words that have been added by retimer 700 and
20 retimer 710. Consequently, a SKIP ordered set that started with three SKIP words now has a total of five SKIP words. Conversely, if an advance is needed to compensate for delayed clock timing, the retimers 700, 710 may remove an existing SKIP word to advance the arrival of subsequent data at the receiving end of the link. SKIP ordered set 730 shows an example of this scenario. SKIP

ordered set 730 contains only one SKIP word as a result of the removal of one SKIP word each by retimer 700 and retimer 710. By compensating for clock tolerances, the link and the Ports on either end of the link can operate in a common clock domain.

In the preferred embodiment, the SKIP word name is K28.0 and the associated 10-bit word is 001111 01000 for negative running disparity and 110000 1011 for positive running disparity. As is the case with the COM and lane identifier words, the SKIP word is insensitive to bit inversion. Other code words exhibiting the same property will also work in alternative embodiments.

Figure 8 shows a block diagram of an adapter 800 configured to convert signals transmitted to and received from a physical link 820. The adapter may be coupled to or otherwise form a part of a port and/or a channel adapter. The adapter 800 is coupled to differential wires or traces 810 in the physical link. Differential signals received from the physical link 820 are detected by a lane receiver 830 that converts the differential signals to a bit stream that is sent to a 10B/8B decoder 850. The decoder converts the 10 bit words received from the individual lanes into 8 bit words that are directed to the FIFO buffers 870. In an alternative embodiment, the FIFO buffers 870 may precede the 10B/8B decoders. After the 10B/8B decoders and FIFO buffers, the 8-bit words are synchronously clocked into a multiplexer or other suitable logic device 880 to reconstruct a single byte stream from the individual byte streams. The byte stream is then sent to a local interface 805 for transmission to the local device 815.

The adapter 800 may also convert signals for transmission to a physical link 820. A byte stream from a local device 815 is detected and transmitted to a demultiplexer 890 that stripes bytes from the single byte stream across a number of individual byte streams. Figure 8 depicts four lanes in the physical link, but this quantity may be different and may depend on whether the link is

coupled to a single channel adapter. The individual byte streams are then coded by the 8B/10B encoders and the resulting bit streams are delivered to lane adapters 840 which convert the bit streams to differential signals for transmission across wire pairs or traces 810 in the physical link 820.

5 The above discussion is meant to be illustrative of the principles and various embodiments of the present invention. Numerous variations and modifications will become apparent to those skilled in the art once the above disclosure is fully appreciated. For example, a physical link with the above properties and characteristics may be constructed with eight or sixteen lanes per link and still operate within the scope of this description. It is intended that the following claims be
10 interpreted to embrace all such variations and modifications.

CLAIMS

What is claimed is:

1. An adapter that comprises:

a lane receiver configured to receive a differential signal on a differential pair of conductors, and configured to convert the differential signal into a sequence of code symbols;

a decoder configured to decode the sequence of code symbols to produce a sequence of received symbols; and

a circuit configured to determine if the sequence of received symbols is incorrect due to inversion of the differential signal, wherein the lane receiver is configured to correct for inversion of the differential signal if the circuit determines inversion exists.

2. The adapter of claim 1, wherein the decoder decodes code symbols from a running-disparity code having a positive running disparity symbol and a negative running disparity symbol for each input symbol.

3. The adapter of claim 2, wherein the circuit is configured to examine a decoded training sequence having a start symbol and a training symbol to determine if inversion exists.

4. The adapter of claim 3, wherein the start symbol has a positive running disparity symbol that is the inverse of the negative running disparity symbol for the start symbol.

1 5. The adapter of claim 3, wherein the inverse of the positive running disparity symbol for the
2 training symbol decodes to a symbol different than the training symbol, and wherein the inverse of
3 the negative running disparity symbol decodes to a symbol different than the training symbol.

1 6. The adapter of claim 5, wherein the circuit locates the start symbol, identifies the training
2 symbol relative to the start symbol, determines if the training symbol has an incorrect value
3 corresponding to the inverse of the positive or negative running disparity symbol for the training
4 symbol.

1 7. The adapter of claim 6, wherein the circuit toggles correction of differential signal
2 inversion if the training symbol has an incorrect value corresponding to the inverse of the positive
3 or negative running disparity symbol for the training symbol.

1 8. The adapter of claim 1, wherein the decoder is an 8B/10B code decoder.

1 9. The adapter of claim 1, further comprising:
2 a second lane receiver configured to receive a second differential signal on a second
3 differential pair of conductors, and configured to convert the second differential
4 signal into a second sequence of code symbols; and
5 a second decoder configured to decode the second sequence of code symbols to produce a
6 second sequence of received symbols,
7 wherein the circuit is further configured to determine if the second sequence of received
8 symbols is incorrect due to inversion of the second differential signal, wherein the

9 second lane receiver is configured to correct for inversion of the second differential
10 signal if the circuit determines inversion exists.

1 10. The adapter of claim 9, wherein the circuit is configured to combine the sequences of
2 received symbols to produce an output symbol stream.

1 11. A method of correcting for differential signal inversion, wherein the method comprises:
2 converting a differential signal into a sequence of code symbols;
3 decoding the sequence of code symbols to form a sequence of received symbols; and
4 determining if the sequence of received symbols is incorrect due to inversion of the
5 differential signal.

1 12. The method of claim 11, further comprising:
2 inverting the sequence of code symbols if inversion is determined.

1 13. The method of claim 11, wherein the code symbols are determined according to a running
2 disparity code having a positive running disparity symbol and a negative running disparity symbol
3 corresponding to each input symbol.

1 14. The method of claim 13, and wherein the determining includes:
2 identifying a training symbol sequence in the sequence of received symbols, said training
3 symbol sequence having a start symbol and a training symbol.

1 15. The method of claim 14, wherein the start symbol has a positive running disparity code
2 symbol that is the inverse of the negative running disparity code symbol for the start symbol, and
3 wherein the inverse of the positive and negative running disparity symbols for the training symbol
4 decode to one or more incorrect symbols different from the training symbol.

1 16. The method of claim 15, wherein the determining further includes:
2 determining if the training symbol has been decoded as one of said incorrect symbols.

1 17. The method of claim 13, wherein said running disparity code is an 8B/10B code.

1 18. The method of claim 11, further comprising:
2 converting a second differential signal into a second sequence of code symbols;
3 decoding the second sequence of code symbols to form a second sequence of received
4 symbols; and
5 determining if the second sequence of received symbols is incorrect due to inversion of the
6 second differential signal.

1 19. The method of claim 18, further comprising:
2 inverting sequences of code symbols for which the sequence of received symbols is
3 determined incorrect due to inversion; and
4 combining the sequences of received signals to form a single output symbol stream.

ABSTRACT

An adapter configured to automatically detect and compensate for differential signal inversion is herein disclosed. In one embodiment, the adapter is part of a computer network having differential conductor pairs conveying differential signals between network devices. The network devices include adapters coupled to transmit and receive signals via the differential conductor pairs. The adapter preferably includes a lane receiver, a decoder, and a synchronization circuit. The lane receiver is configured to receive a single differential signal and to convert the differential signal into a sequence of code symbols. The decoder decodes the code symbols to produce a sequence of received symbols. The synchronization circuit examines the sequence of received symbols to determine if it is incorrect due to inversion of the differential signal, and if so, it causes the lane receiver to correct for the differential signal inversion. It is expected that the received symbol sequence will include a training symbol sequence which will have a start symbol whose decoded value is unaffected by differential signal inversion, and a training symbol whose decoded value is indicative of the presence or absence of inversion. The synchronization circuit, upon identifying the training sequence, will thus be able to determine whether inversion exists and be able to automatically correct for it.




FIG. 1




FIG. 2




FIG. 3

Physical Link Lane Identifiers

LANE IDENTIFIER	8B/10B CODE NAME	NEGATIVE RD	POSITIVE RD
LID 0	D0.0	10011 10100	01100 01011
LID 1	D1.0	01110 10100	10001 01011
LID 2	D2.0	10110 10100	01001 01011
LID 3	D4.0	11010 10100	00101 01011
LID 4	D8.0	11100 10100	00011 01011
LID 5	D15.0	01011 10100	10100 01011
LID 6	D16.0	01101 10100	10010 01011
LID 7	D23.0	11101 00100	00010 11011
LID 8	D24.0	11001 10100	00110 01011
LID 9	D27.0	11011 00100	00100 11011
LID 10	D29.0	10111 00100	01000 11011
LID 11	D30.0	01111 00100	10000 11011

FIG. 6




FIG. 7

Express Mail Label No. JI497841418US

Attorney Docket No. 1662-28800

Client Docket No. P00-2998

DECLARATION

SOLE/JOINT INVENTOR
ORIGINAL/SUBSTITUTE/CRP

As a below named inventor, I hereby declare that: my residence, post office address, and citizenship are as stated below next to my name. I believe I am the original, first, and sole inventor (if only one name is listed below) or a joint inventor (if plural inventors are listed below) of the subject matter which is claimed and for which a patent is sought on the invention entitled: **HIGH-SPEED INTERCONNECTION ADAPTER HAVING AUTOMATED CROSSED DIFFERENTIAL PAIR CORRECTION**, as described in the specification attached.

I hereby state that I have reviewed and understand the contents of the above identified specification, including the claims, as amended by any amendment referred to above; that I do not know and do not believe the same was ever known or used in the United States of America before my or our invention thereof, or patented or described in any printed publication in any country before my or our invention thereof or more than one year prior to this application; that the invention has not been patented or made the subject of an inventor's certificate issued before the date of this application in any country foreign to the United States of America on an application filed by me or my legal representative or assigns more than twelve months prior to this application; and that I acknowledge the duty to disclose information of which I am aware which is material to the examination of this application in accordance with Title 37, Code of Federal Regulations § 1.56(a). Such information is material when it is not cumulative to information already of record or being made of record in the application, and

- (1) it establishes, by itself or in combination with other information, a prima facie case of unpatentability of a claim; or
 (2) it refutes, or is inconsistent with, a position the applicant has taken or may take in:
- (i) opposing an argument of unpatentability relied on by the Office, or
 (ii) asserting an argument of patentability.

I hereby claim foreign priority benefits under Title 35, United States Code § 119 of any foreign application(s) for patent or inventor's certificates listed below and have also identified below any foreign application(s) having a filing date before that of the application(s) on which priority is claimed:

COUNTRY	APPLICATION NUMBER	DATE OF FILING	PRIORITY CLAIMED UNDER 35 USC 119
			<input type="checkbox"/> YES <input type="checkbox"/> NO

I hereby claim the benefit under Title 35 United States Code § 120 of any United States application(s) listed below and, insofar as any subject matter of any claim of this application is not disclosed in the prior United States Application, I acknowledge the duty to disclose material information as defined in Title 37, Code of Federal Regulations § 1.56(a) which occurred between the filing date of the prior application and the national PCT international filing date of this application:

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code and that such willful false statements may jeopardize the validity of the application or any patent issued thereon.

FULL NAME OF SOLE OR FIRST INVENTOR		INVENTOR'S SIGNATURE	DATE
William P. BUNTON		<i>William P. Bunton</i>	6-19-00
RESIDENCE		CITIZENSHIP	
415 Greenway Drive, Pflugerville, Texas, 78660		U.S.A.	
POST OFFICE ADDRESS			
SAME AS ABOVE			
FULL NAME OF SECOND JOINT INVENTOR		INVENTOR'S SIGNATURE	DATE
John KRAUSE			
RESIDENCE		CITIZENSHIP	
1310 E. University Avenue, Georgetown, Texas, 78626		U.S.A.	
POST OFFICE ADDRESS			
SAME AS ABOVE			
FULL NAME OF THIRD JOINT INVENTOR		INVENTOR'S SIGNATURE	DATE
Patricia L. WHITESIDE		<i>Patricia L. Whiteside</i>	6-19-00
RESIDENCE		CITIZENSHIP	
15212 Quiet Pond Court, Austin, Texas, 78728		U.S.A.	
POST OFFICE ADDRESS			
SAME AS ABOVE			